AlzFed-XAI: Privacy-Preserving Alzheimer's Diagnosis

High-Fidelity, Interpretable AI using Federated Learning

Authors: Md. Abdur Rahman, Md. Tofael Ahmed Bhuiyan, Abdul Kadar Muhammad Masum

Department of Computer Science and Engineering, Southeast University, Bangladesh

Paper ID: 56 (Track 2)

2nd December, 2025

Introduction & Motivation

- **The Global Crisis:** Alzheimer's Disease (AD) is a leading cause of dementia. Early detection via MRI is crucial for effective patient management.
- **Deep Learning Promise:** Convolutional Neural Networks (CNNs) excel at identifying pathological changes in scans, but they require massive, diverse datasets to generalize well.
- **The Bottleneck:** Strict privacy regulations (HIPAA, GDPR) make centralized data collection illegal or impractical, creating "Data Silos" that hinder AI progress.
- **Our Objective:** To develop AlzFed-XAI, a framework that enables collaborative learning across institutions without sharing patient data, while ensuring the results are explainable to clinicians.

The Privacy Challenge in Medical AI

The Conflict

Deep Learning needs massive data, but privacy laws (HIPAA/GDPR) lock patient data in isolated hospitals ("Data Silos").

The Solution

Federated Learning. We bring the model to the data, not the data to the model.

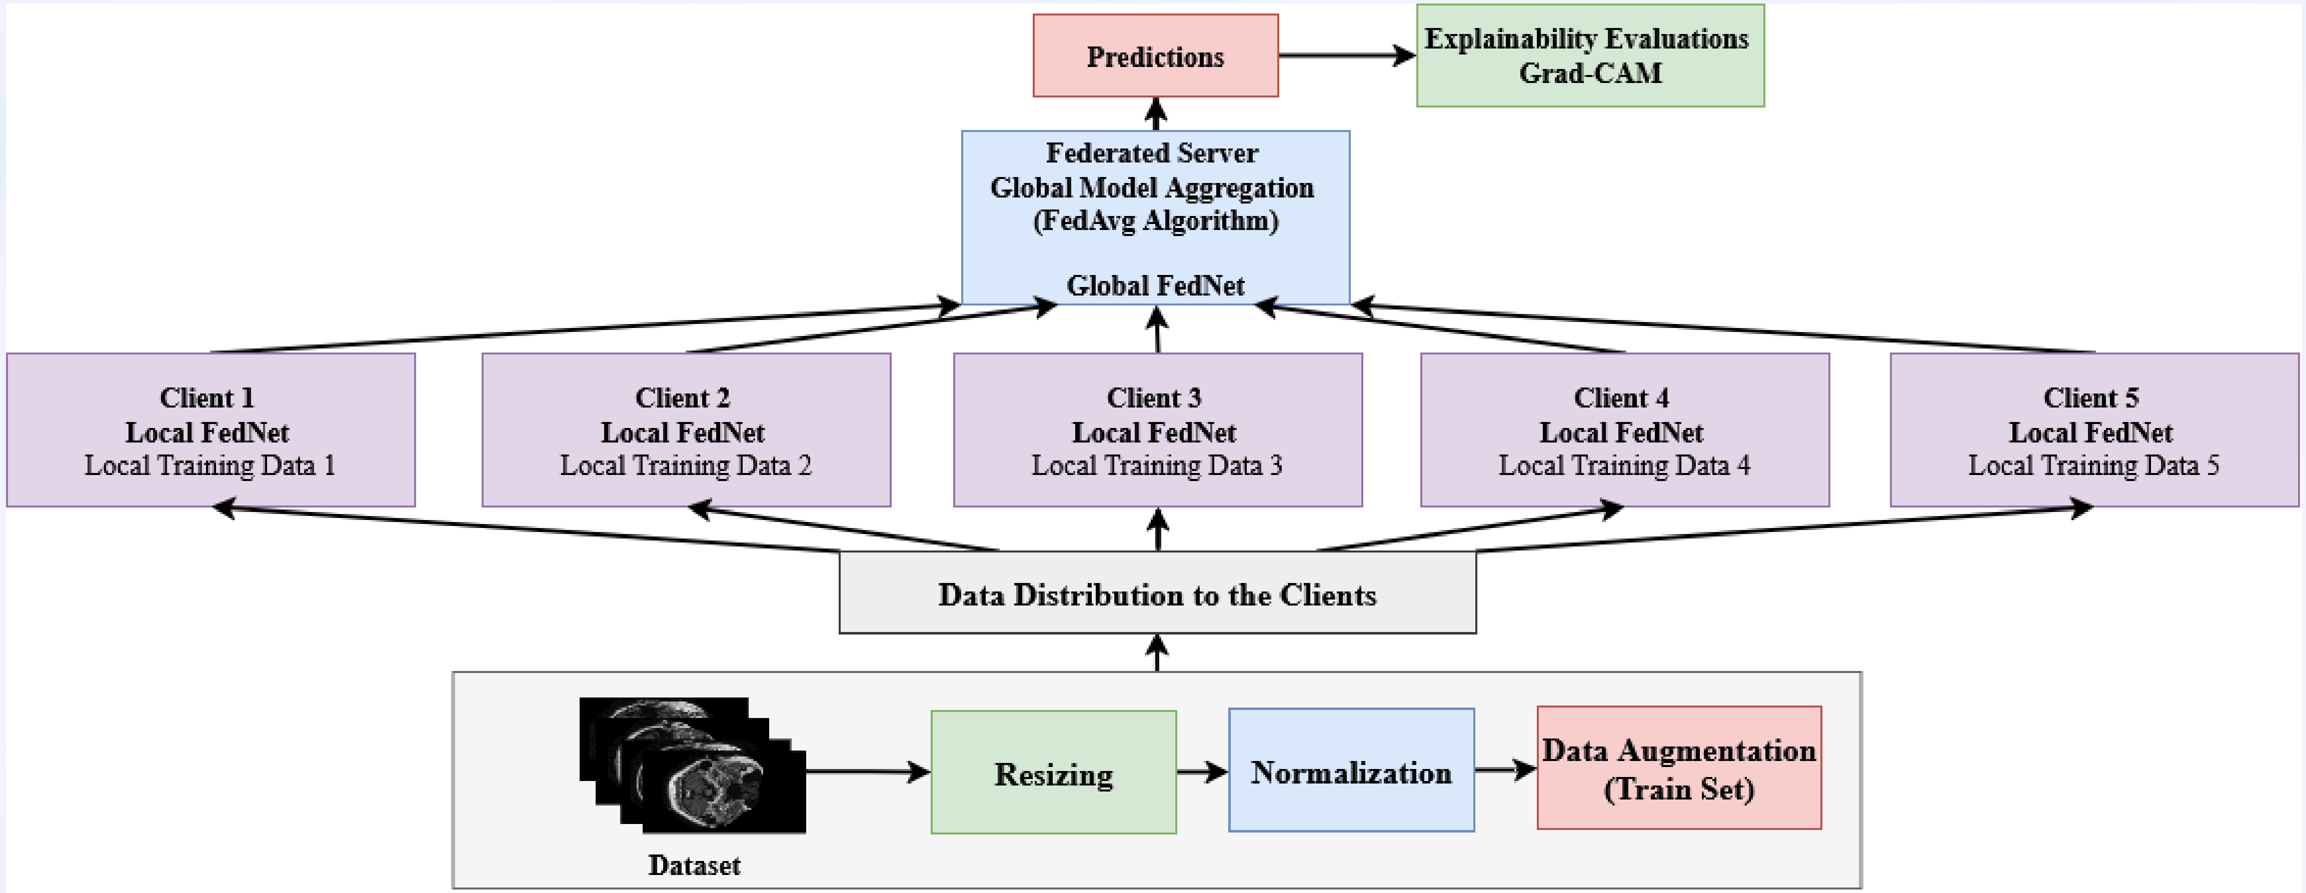

Figure: Proposed AlzNet-XAI Framework

2nd December, 2025

FedNet: Engineered for Efficiency

- **The Model:** A custom Lightweight CNN designed for edge devices.
- **Key Stats:** Only **378,780 parameters** (vs. millions in standard models).
- **Tech Stack:** Uses Mobile Inverted Bottleneck (MBConv) and Depthwise Separable Convolutions for speed without losing accuracy.

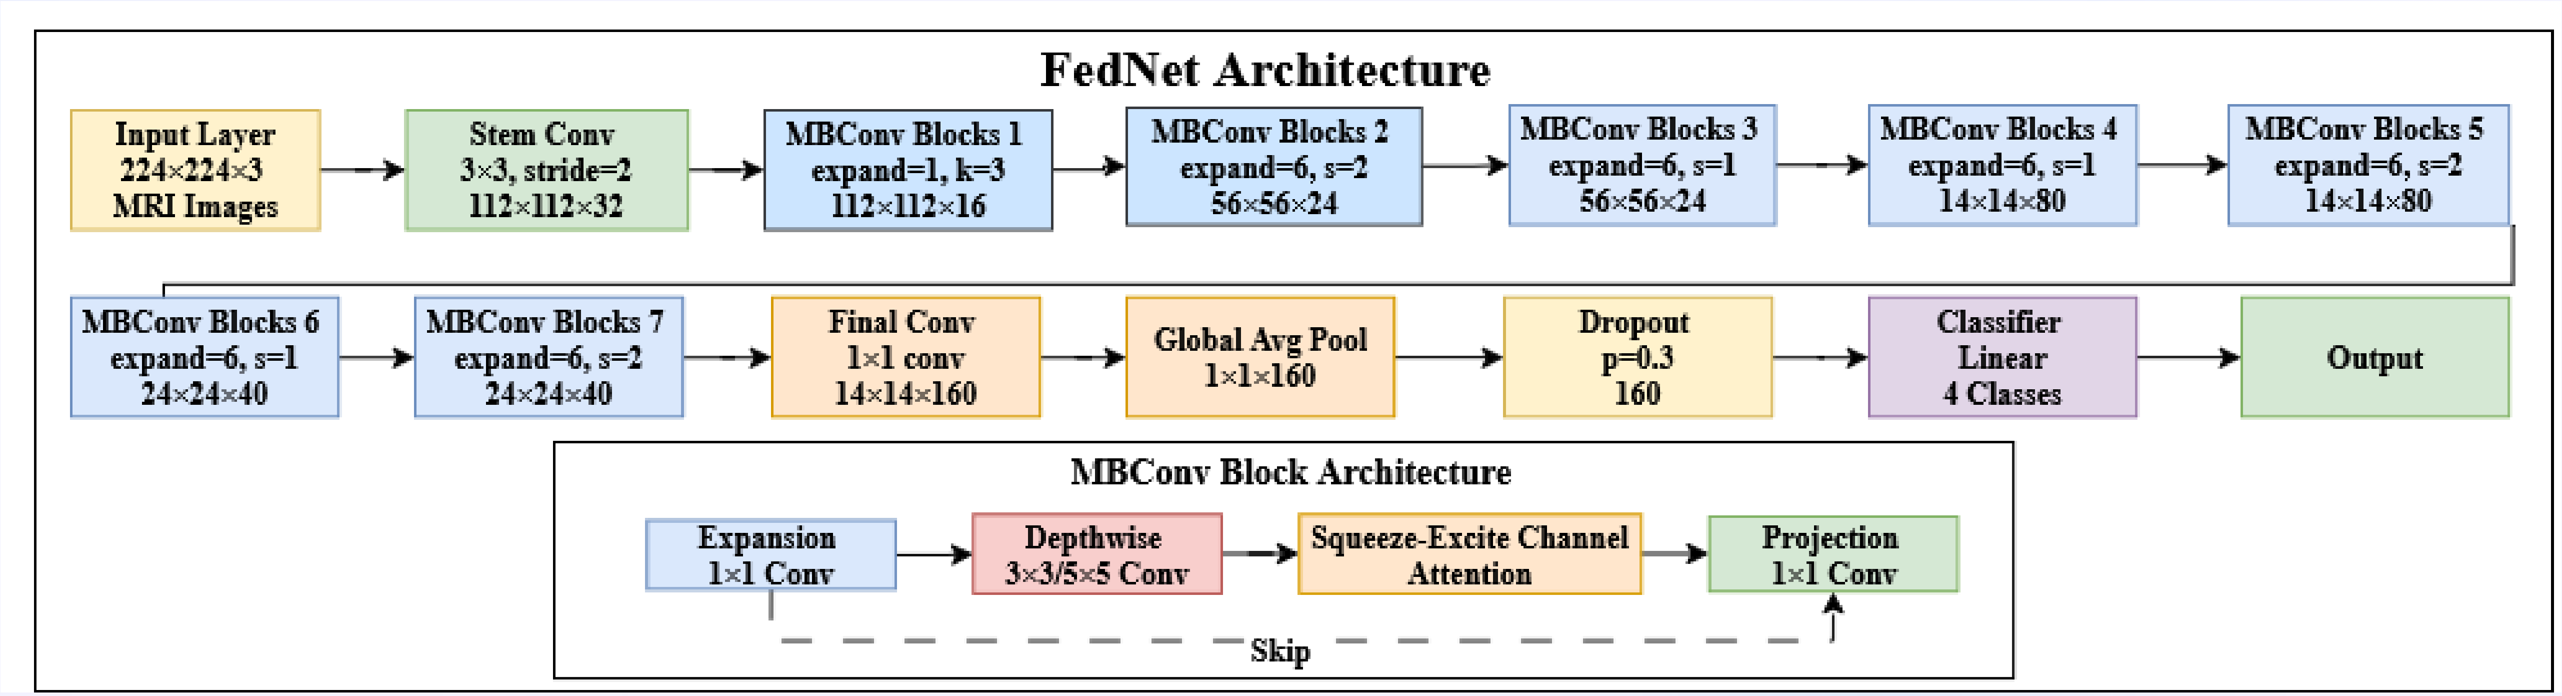

Figure: The architecture block diagram showing the flow from input MRI to the MBConv blocks and final classification head.

Federated Optimization Protocol

- **Algorithm:** Federated Averaging (FedAvg).
- **Process:**

01

Server broadcasts global weights.

02

Clients train locally for 3 epochs.

03

Server aggregates updates to improve the global model.

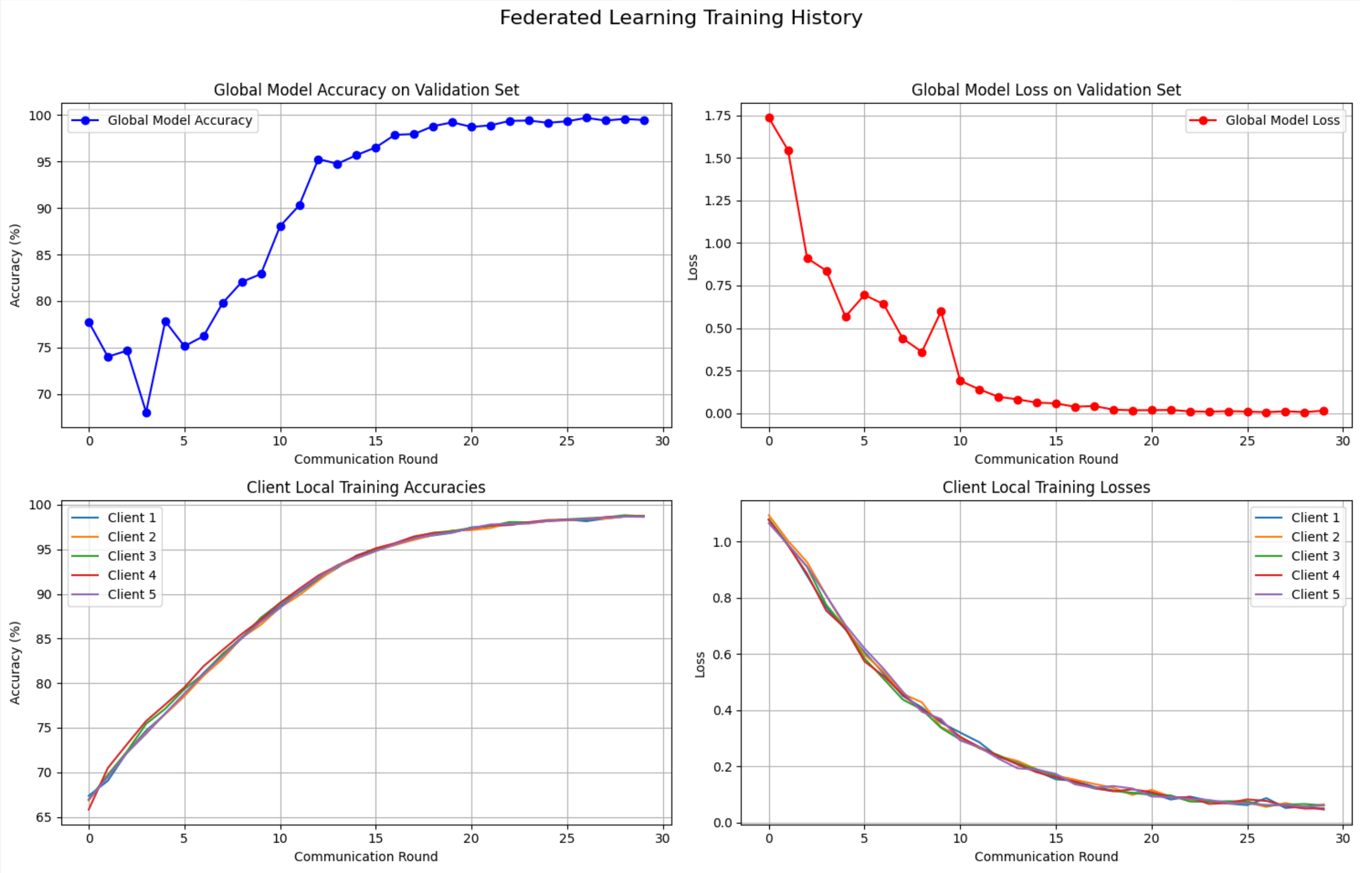

Figure: Training dynamics charts. The top line shows Global Accuracy rising to 99%, while the bottom lines show local client losses decreasing smoothly.

2nd December, 2025

Experimental Setup & Data

- **Dataset:** OASIS-1 MRI (Cross-sectional).
- **Simulation:** 5 distinct clients with non-overlapping data partitions.
- **Environment:** NVIDIA Tesla P100 GPU.

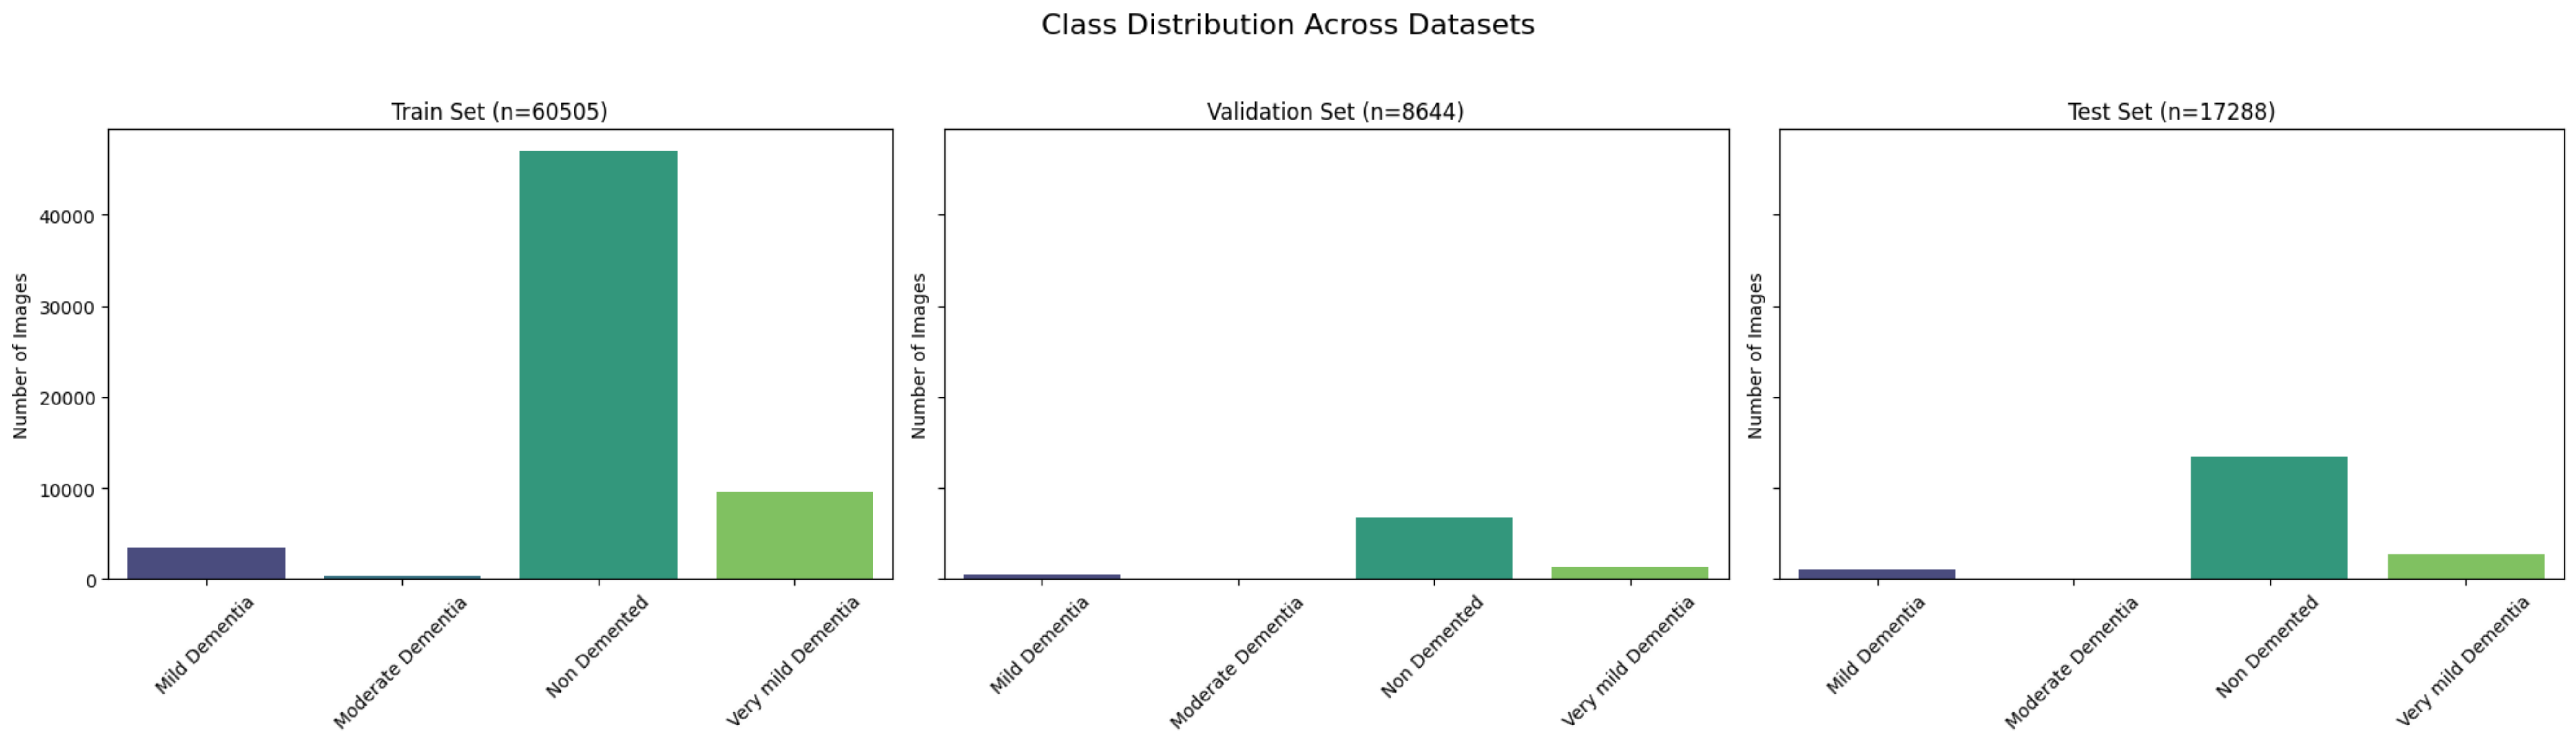

Figure: Bar charts showing the distribution of images across the 4 classes (Non-Demented, Very Mild, Mild, Moderate) for Train, Validation, and Test sets.

Near-Centralized Performance

Table 1: Performance comparison of FedNet baseline and proposed AlzFed-XAI framework.

Model	Test accuracy (%)	Precision (macro)	Recall (macro)	F1-score (macro)
FedNet	99.9364	0.9980	0.9997	0.9988
AlzFed-XAI	**99.7281**	**0.9959**	**0.9982**	**0.9970**

- **Impact:** We achieved high privacy with a negligible performance drop (< 0.2%).
- **Robustness:** The model handles class imbalance perfectly with a 0.997 F1-Score.

Trust via Explainability (Grad-CAM)

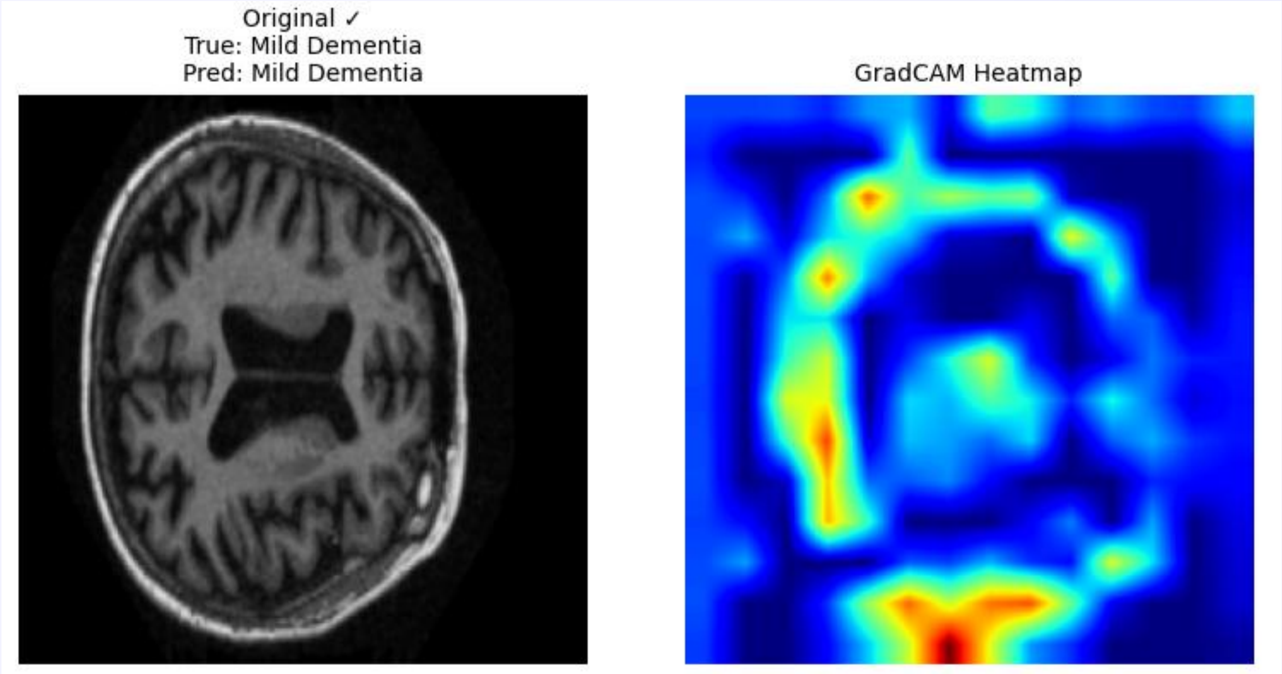

- **The "Black Box" Problem:** Doctors need to know *why* an AI made a diagnosis.
- **Our Solution:** Gradient-weighted Class Activation Mapping (Grad-CAM).
- **Clinical Validation:** The model focuses on the **Temporal and Parietal lobes** regions known to atrophy in Alzheimer's proving it learns biology, not noise.

Figure: Brain MRI scans overlaid with heatmaps. The red "hot spots" show exactly which part of the brain the model looked at to detect dementia.

Conclusion & Future Scope

Viability

Federated Learning is ready for sensitive medical diagnostics.

Efficiency

High accuracy is possible on low-resource hardware using FedNet.

Trust

Interpretability (XAI) is the key to clinical adoption.

Future:

Deploying on real-world, non-IID data across different hospital scanner types.

Thank you

2nd December, 2025